Making CDN and ISP Routings Symbiotic

Varun Khare and Beichuan Zhang
Department of Computer Science, University of Arizona
Email: {vkhare, bzhang}@cs.arizona.edu

***Abstract*—Internet Service Providers (ISPs) route traffic at the IP layer with the preference of less inter-carrier payments while Content Distribution Networks (CDNs) route traffic at the application layer with the preference of better application performance. Such mismatch of routing preferences leads to conflicts that eventually result in higher operational cost for both ISPs and CDNs. In this paper, we propose to make CDN and ISP routing mutually beneficial through ISP's non-uniform bandwidth charging and CDN's bandwidth cost-aware request routing. More specifically, ISPs charge different prices for traffic that traverses different types of inter-domain links and CDNs, in routing user requests to their servers, try to minimize their ISP payments by taking the pricing information into consideration. We evaluate the solution in large scale simulations. The greedy solution presents the lowest bandwidth cost for CDNs but at the expense of network performance for users. With end-to-end delay introduced as a constraint in the optimization process, the solution maintains good network performance for users while achieving significant savings in bandwidth cost. Compared with conventional nearest-available policy in CDN request routing, our solution moves significant amount of inter-domain traffic from provider routes to peer or customer routes, reducing operational costs for ISPs and CDNs.**

I. Introduction

Internet Service Providers (ISPs) interconnect amongst each other based on commercial agreements that determine the flow of money between them and thereby influence the flow of traffic between them as well. Typical commercial agreements are that customer networks pay provider networks for data transit services and peer networks do not pay each other explicitly [14]. ISPs then administer routing policies over inter-domain traffic to minimize their inter-carrier payments. In typical policies, ISPs prefer to route traffic through customers rather than peers to gain revenue from customers, and through peers rather than providers to avoid incurring provider charges [24].

Content Distribution Networks (CDN) route content traffic at the application layer to provide better application performance [19]. CDNs are overlay networks composed of a large set of dedicated servers that are deployed in many ISPs. When a user requests certain content, a CDN redirects the request to a server that has the content. Then the ISP in which the server resides will route the requested content traffic to the ISP of the user. At present, CDN's request routing mechanism only considers metrics of application performance, *e.g.* network delay [12] and server load [11], but not underlying ISPs' routing policies.

The mismatch of routing preference between CDNs and ISPs can result in higher operational costs for both. CDN's request routing determines the source (*i.e.*, the server) and the destination (*i.e.*, the user) of the content traffic, which then will be routed from the server to the user through underlying ISP networks. There can be a conflict between CDNs and ISPs in this process. For instance, between two servers located in two different ISPs, CDNs may choose the server that has a shorter end-to-end delay to the user for better application performance, but ISPs may prefer the other server whose traffic will take customer route rather than provider route to save on the inter-carrier payment. With increasing commercial success of CDNs in delivering contents at the Internet scale [1], [6], the problem of mismatched routing preferences between CDNs and ISPs is becoming more prominent.

Actually CDNs have the economic incentives to resolve such preference mismatch but do not have necessary information or mechanism to pursue such actions. CDNs pay the ISPs for using the underlying networks based on the traffic volume that CDN servers send and receive. If an ISP's operational cost goes up due to the increase of inter-carrier payments, it will eventually be reflected in higher ISP charges to CDNs. Therefore CDNs and ISPs have the common interest in reducing ISPs' inter-carrier payments, which will reduce the operational costs for both of them. The main question is that, for CDNs to cooperate, what information needs to be shared between ISPs and CDNs, and through what mechanisms the cooperation is to be done. Ideally the shared information should be minimal and the cooperation should not introduce any unnecessary dependency between ISPs and CDNs.

We propose that in order to economically incentivize CDNs to consider underlying ISPs' routing preference, ISPs should charge differently for content traffic depending upon the type of inter-domain route it takes and make this pricing information available to CDNs. For example, content traffic taking provider route is more expensive than content traffic taking peer route. Such non-uniform pricing on content traffic does not change frequently since the underlying inter-ISP relationship is relatively stable, thus it can be readily shared with CDNs. ISPs already monitor their inter-domain traffic at border routers and calculate their bills based on the traffic volume. The only extra thing they need to do is to use different prices for different types of inter-domain traffic generated by CDN servers.

We devise a novel CDN request-routing mechanism, called COst aware REquest routing for Overlay Multicast Networks (CORE-OMN), whose objective is to minimize CDN payment to ISPs. We formulate the user-request routing problem, with given ISP routes and corresponding prices, as an optimization problem where each user to server assignment contributes an ISP charge. We outline the optimal solution but focus upon a fast greedy heuristic that assigns new user request to the server that incurs the least marginal ISP charge. CORE-OMN by minimizing ISP payments alone may negatively impact application performance, and thus metrics of application performance must be included as constraints to the optimization problem. In this paper we use end-to-end delay as such a constraint, which turns out to be effective. With ISP's routing preference expressed economically and factored into CDN's request-routing, more content traffic will take cheaper IP routes within the delay constraint, which reduces operational cost for ISPs and CDNs.

The principles of the solution are applicable to all the services offered by CDN. As the first step in this direction, this paper focuses on the overlay multicast service only, which involves huge amount of traffic with significant impacts on network operations. CDNs such as Akamai [2] and Limelight [4] have been the most successful in delivering streaming video content to millions of viewers Internet wide, including popular political events [6] and live sports events [1]. Recent studies [7] also predict dramatic increase in the demand for online streaming video content, which means that multicast CDNs in the future will be delivering a much larger quantity of such inter-domain traffic thereby increasing their impacts upon the operational costs of underlying ISPs.

In evaluation, we use simulations to compare variations of CORE-OMN against nearest-available request routing policy that is commercially adopted in many CDNs [12]. CORE-OMN-Greedy achieves the lowest bandwidth cost but at the expense of network performance for end users. CORE-OMN-Delay provides good network performance for end users with significant savings in ISP payment. Compared with nearest-available request-routing, CORE-OMN moves significant amount of inter-domain traffic from provider routes to either peer or customer routes, reducing operational costs for both CDNs and ISPs. We also design and evaluate a distributed mechanism to maintain low bandwidth cost as group membership changes with user join and leave.

The rest of the paper is organized as follows: Section II gives brief background on multicast CDN and ISP charging models. Section III presents design overview. Section IV formulates CORE-OMN user assignment problem and proposes solutions. Section V presents CORE-OMN protocol along with user movement problem and its solution. Section VI presents evaluation methodology and simulation results. We discuss related work in Section VII and conclude the paper in Section VIII.

II. ISP Charge and Multicast CDN Network Model

ISPs charge business customers based on the traffic volume that customers send and receive within a charging period, *e.g.*, one month. There are two popular ways to determine the charging volume: 95-percentile and total volume. In the former approach, traffic volume for each 5-minute is recorded for a month and the 95-percentile over the sorted values is used as charging volume. In the latter approach, the total traffic volume within a month is the charging volume. ISPs determine the final bill using a charging function with the computed charging volume. The charging function, as reported [27], [13], follows

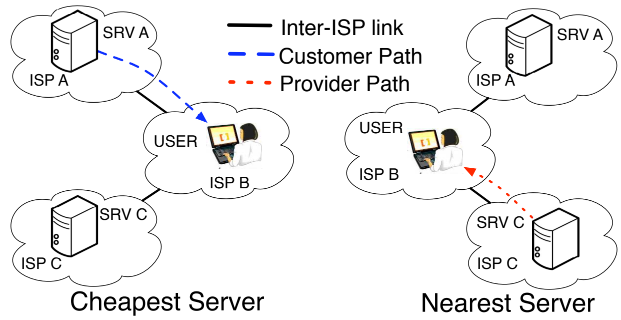

Fig. 1. Impact of user assignment policy on inter-domain traffic

economy of scale i.e. the unit price per Mbps drops as the total purchased bandwidth goes up.

A multicast CDN consists of a content delivery infrastructure, a request-routing mechanism, and an overlay multicast distribution mechanism [20]. The delivery infrastructure consists of dedicated servers with reserved bandwidth, deployed at strategic locations within the Internet such as Point of Presence (PoP) of various ISPs [12]. The distribution mechanism is essentially an overlay multicast protocol that connects the servers into one or multiple dissemination trees through which the data will be delivered from the source server to edge servers. Many protocols have been designed to build such dissemination trees, *e.g.*, OMNI [8], AMCast [21], HMTP [25] and Narada [9], each optimizing particular application performance metrics.

The focus of this paper is the request-routing mechanism, which redirects user requests to servers with requested content and available bandwidth subject to user assignment policies. The conventional user assignment policy is "nearest-available", which selects the server that is nearest to user and has enough available bandwidth to send another stream of data to user [12]. There are variants of the nearest-available policy that consider other factors such as server load [11], cache locality [26], and robustness [23]; however they all focus on application performance and not the impact upon ISP charges.

The user assignment policy affects the operational costs of both CDNs and ISPs. Consider the hypothetical network topology in Figure 1, where two servers SRV_A and SRV_C are able to serve the same content to user. ISP_A is provider of ISP_B, and ISP_C is customer of ISP_B. Also, SRV_A has longer network delay to user than SRV_C. The nearest-available policy would choose SRV_C to serve user, which means ISP_C will have to pay its provider for the traffic incurred, and the cost will eventually be transferred to the CDN too. But if the CDN tries to minimize its ISP payment, it would choose SRV_A to serve user since ISP_A does not pay its customers, and the reduced cost to ISP_A will eventually benefit the CDN too. In this paper we generalize the idea in this simple example to achieve mutually beneficial cooperation between CDNs and ISPs.

III. Design Overview

This section outlines the enabling mechanism for CDN and ISP cooperation.

ISP Perspective: ISPs, at present, charge purely based on traffic volume, which offers no incentive for CDNs to consider underlying ISPs' routing preference. Traffic of the same amount costs differently to an ISP depending upon the type of link the traffic traverses. In general ISPs prefer the traffic to stay within their own network, *i.e.*, so-called "on-net" traffic. Thereafter ISPs regulate their "off-net" traffic, which leaves their networks, through inter-domain routing policies. None of this is reflected in ISP charges to CDNs, leaving CDNs clueless even if they are willing to cooperate.

ISPs need to charge different prices for on-net traffic and off-net traffic generated by CDN servers to recover its operational costs and more importantly, provide economic incentives for CDN to consider ISPs' routing preferences. A natural setting is to model prices after their respective transmission costs, *e.g.*, ISPs can set charges for provider, peer, customer off-net traffic and on-net traffic in that order from most expensive to least expensive where each charging function still follows economy of scale. To calculate CDNs' final bills, ISPs need to measure the traffic volume on each of its inter-domain links, which is already commonly done in commercial ISPs by using tools like NetFlow. The ISPs need to identify traffic of CDN servers by source and destination IP addresses, and apply different prices to different types of traffic. Such non-uniform pricing information needs to be shared with CDNs. It can be treated in the same way as today's ISP charging information, which is usually part of the contract with the CDNs. Therefore the change to ISP operations is small.

CDN Perspective: Under the non-uniform ISP charging, CDNs are incentivized to minimize their ISP payments by adjusting their request routing policy, which is examined in detail in the next section. In order to do this, CDNs need to know that given a destination address (*i.e.* IP address of user), what routes are available and their associated prices from the server sites. Note that CDNs do not need to know ISP's inter-domain relationship or why ISPs charge a particular price on a particular route. All CDNs need to know is what will be the ISP charge if traffic is sent from a server to a particular IP.

One way to obtain such routing and pricing information is through a passive BGP session, from which the ISP routers announce their routing updates augmented with pricing information in the form of BGP community attribute PATH_PRICE, and CDN's control servers receive this information and make request-routing decisions. A CDN's control server can set up such passive BGP sessions with all the ISPs where CDN has presence, collect necessary pricing information, and carry out the optimization process to determine which server in which ISP should serve the user. This approach of passive BGP session is compatible with current operations and incurs minimal cost to both CDNs and ISPs.

Now CDNs and ISPs share the minimal information needed, and for most part still operate independently. But the result will make economic sense for both of them as the cooperation is done implicitly through CDNs' cost-aware request-routing.

TABLE I
NOTATION FOR CORE-OMN USER ASSIGNMENT PROBLEM

K $\text{SRV}_{k,j}$	Number of servers deployed within ISPs. Server SRV_k deployed in ISP_j with bandwidth B.
N $\text{U}_{i,j}$	Number of interested users in group. Interested user U_i in ISP_j consuming b bandwidth
C_j^{on}	ISP_j charging function for on-net traffic.
C_j^{off}	ISP_j coarse charging function for off-net traffic.
C_j^{prov} C_j^{peer} C_j^{cust}	ISP_j charging function for off-net traffic being sent over its provider, peer and customer link respectively.
U_k^{on}	Number of on-net users assigned to SRV_k.
U_k^{off}	Number of off-net users assigned to SRV_k.
U_k^{prov} U_k^{peer} U_k^{cust}	Number of off-net users assigned to $\text{SRV}_{k,j}$ where corresponding off-net traffic from ISP_j is sent on provider, peer and customer link respectively.
$\text{dist}_{i,k}$	Overall delay experienced by user U_i at SRV_k

IV. CORE-OMN REQUEST ROUTING

In this section, we formulate the user assignment problem based on the modified ISP charging model. From henceforth we use bandwidth cost and ISP charge interchangeably.

A. General CORE-OMN User Assignment Problem

We introduce the notation in Table I and state the general CORE-OMN user assignment problem formally as: Given: (1) K servers s.t. any $\text{SRV}_{k,j}$ with bandwidth B deployed in ISP_j uses C_j^{on} and C_j^{off} charging function for on-net and off-net traffic [1] respectively; (2) N interested users s.t. $\text{U}_{i,j}$ present in ISP_j consumes b bandwidth. Find the user assignment which minimizes bandwidth cost of the multicast group $\sum_{k=1}^{K} C_j^{on}(U_k^{on} \cdot b) + C_j^{off}(U_k^{off} \cdot b)$ under the following constraints: (1) all users are assigned: $\sum_{k=1}^{K} U_k^{on} + U_k^{off} = N$ and (2) bandwidth constraint at each server is met: $(U_k^{on} + U_k^{off}) \cdot b \leq B$. Bandwidth consumption in supporting multicast group occurs in the overlay tree from origin to edge servers and from edge server to users. But the bandwidth consumed by users is orders of magnitude greater as users are far more than servers in the overlay tree, and therefore bandwidth consumed by users is only considered in aforementioned formulation.

Offline Dynamic Programming Solution: Dynamic programming approach provides optimal solution to the CORE-OMN user assignment problem since the optimal solution to distribute N users amongst K servers contains within it the optimal solution to the sub-problem of distributing n users amongst k servers where n≤N and k≤K. Let cost(n, k) be the optimal cost for allocating n users amongst k servers. The evaluation starts at cost(n,1) where n=1,...,N with k=1 forcing every users to be assigned to $\text{SRV}_{k=1}$. While assigning users to SRV_k the respective on-net or off-net bandwidth cost is applied depending upon ISP locality of sever and user. Finally evaluating

[1] In this formulation we consider coarse charging function for off-net traffic without distinguishing traffic sent over provider, peer or customer links.

$$cost(n,k) = \begin{cases} C_1^{on}(\sum_{i=0}^{n} U_{i,ISP=1} \cdot b) + C_1^{off}(\sum_{i=0}^{n} U_{i,ISP\neq 1} \cdot b) & k=1, B_1 \geq n \cdot b \\ \min_{0<i\leq n} cost(n-x,k-1) + C_k^{on}(\sum_{i=0}^{x} U_{i,ISP=k} \cdot b) + C_k^{off}(\sum_{i=0}^{x} U_{i,ISP\neq k} \cdot b) & k>1 \end{cases}$$

cost(N,K) gives the optimized cost of deploying the multicast group. Tracking the user distribution of the sub-problems allows us to find the user distribution for cost(N,K), the solution to the user assignment problem. The runtime of the algorithm is O(K.N^2) and the space complexity is O(K.N). The algorithm considers all the sub-problems in the space of O(K.N) therefore for any change in membership of the multicast group, say ΔN, requires O(K$\cdot\Delta$N^2) computation to evaluate the cost-optimal user assignment. Since the dynamic programming solution is computationally expensive and slow for assigning users in a flash-crowd there is a need for an online greedy heuristic.

Online CORE-OMN Greedy Heuristic: CORE-OMN-Greedy heuristic, by identifying the type of traffic generated for any server-user assignment and its associated ISP charges is able to perform bandwidth cost-efficient user assignment. The greedy heuristic exploits the fact that charging function for on-net traffic is always cheaper than charging function for off-net traffic. For CDN this means on-net server which is in the same ISP location as user is cheaper and therefore preferred than any off-net server present in different ISP location as user for bandwidth cost-efficient user assignment. On-net servers move content closer to users and by acting as cache site minimize ISP's network access cost. The greedy heuristic preserves these benefits by giving preference to any available on-net server for user assignment. Thereafter greedy heuristic assigns user to off-net server offering the lowest marginal off-net bandwidth cost.

CORE-OMN-Greedy presented in Algorithm 1 performs bandwidth cost-efficient user assignment after classifying servers into on-net and off-net type and thereafter comparing the respective marginal bandwidth cost offered by each server. Servers can be classified as on-net and off-net type for any given user by comparing the ISP location of server against user. Server locations are known as part of the delivery infrastructure and user is assumed to be located within ISP originating the BGP path to user's destination address. Thereafter CORE-OMN-Greedy compares marginal bandwidth cost offered by each server which depends upon server load and ISP charging function. MetaCDN [11], adopted in commercial CDNs such as Akamai [2], describes the collection and maintenance of server loads and the ISP charging function is known as part of the delivery infrastructure.

CORE-OMN-Greedy assigns user to server offering cheapest marginal bandwidth cost. CORE-OMN-Greedy evaluates the locally cheapest on-net server (SRV$_{on}$) and the globally cheapest off-net server (SRV$_{off}$) offering the least respective marginal bandwidth cost from the respective list of servers. In each server list, the server (SRV$_{pos}$) offering least immediate marginal bandwidth cost is evaluated as the candidate cheapest server. However points of intersection, say at p, may exist between charging functions of SRV$_{pos}$ and another server, say SRV$_{int}$, representing the bandwidth cost contention region since c$_{pos}$(p) = c$_{int}$(p) where c$_{pos}$ and c$_{int}$ are ISP charging functions of SRV$_{pos}$ and SRV$_{int}$ respectively. The cheapest server at the intersection point is evaluated by comparing c'_{pos}(p) against c'_{int}(p). After evaluating SRV$_{on}$ and SRV$_{off}$, CORE-OMN-Greedy compares their respective marginal bandwidth cost and assigns user to server offering lowest marginal bandwidth cost. Since on-net charging function is expected to be always cheaper than off-net charging function, the default user assignment is always to the on-net server. The runtime of CORE-OMN-Greedy is O(K) suitable for assigning users in a flash crowd where group membership can change rapidly.

Algorithm 1 Online CORE-OMN-Greedy User Assignment

Greedy User Assignment
SRVList = list of servers with available bandwidth
pUser = user to be assigned consuming **b** bandwidth
for all SRV$_k$ $\in$ SRVList **do**
  if SRV$_k\cdot$ISP$_j$ == pUser$\cdot$ISP **then**
    OnNet-SRVs$\cdot$add(SRV$_k$); Cost$_k$ = C_j^{on}(b)
  else
    OffNet-SRVs$\cdot$add(SRV$_k$); Cost$_k$ = C_j^{off}(b)
SRV$_{off}$ = Cheapest-Off-Net-Server(OffNet-SRVs)
SRV$_{on}$ = Cheapest-On-Net-Server(OnNet-SRVs)
pUser **joins** SRV with **cheapest marginal bandwidth cost**

B. Extended CORE-OMN User Assignment Problem

Extended CORE-OMN user assignment problem takes into account all the different types of off-net traffic which can be generated at a deployed server site. In the general problem formulation, ISPs only distinguish between on-net and off-net traffic which is useful for server sites placed within customer networks and small regional providers where mostly on-net traffic is served and off-net traffic can only take the default provider route. However Tier-1 and Tier-2 networks offer more choice of routes to its off-net traffic due to their interconnections with several ISPs which involve different types of commercial agreements and therefore different types of intercarrier payments.

The Extended CORE-OMN user assignment problem is stated formally as: Given: (1) K servers s.t. any SRV$_{k,j}$ with bandwidth B deployed in ISP$_j$ uses C_j^{on} and C_j^{prov}, C_j^{peer} and C_j^{cust} charging functions for on-net and different types of off-net traffic respectively; (2) N interested users, U$_{i,j}$ consuming b bandwidth and present in ISP$_j$. Find the user assignment which minimizes bandwidth cost of the group $\sum_{k=1}^{K} C_j^{on}(U_k^{on} \cdot b) + C_j^{prov}(U_k^{prov} \cdot b) + C_j^{peer}(U_k^{peer} \cdot b) + C_j^{cust}(U_k^{cust} \cdot b)$ under the following constraints: (1) all users are assigned: $\sum_{k=1}^{K} U_k^{on} + U_k^{prov} + U_k^{peer} + U_k^{cust} = N$ and (2) bandwidth constraint for each server is met: $(U_k^{on} + U_k^{prov} + U_k^{peer} + U_k^{cust}) \cdot b \leq B$.

Online Extended CORE-OMN-Greedy Heuristic: Extended CORE-OMN-Greedy heuristic regulates the off-net traffic generated at server sites in accordance with the underlying ISPs' routing policies due to the economic incentives presented by the ISPs in the form of the modified charging functions. The greedy heuristic reduces the operational ISP charges by transferring CDN servers' off-net traffic from costly provider or peer IP routes to cheaper customer IP routes. For CDN this means customer off-net server where off-net traffic traverses over cheaper customer IP route is preferred for user assignment over provider or peer off-net server where off-net traffic traverses over costly provider or peer IP route respectively. The greedy heuristic still gives preference to on-net server if available, but re-distributes the off-net traffic to reduce operational cost for both CDN and ISPs.

Extended CORE-OMN-Greedy presented in Algorithm 2 considers the different off-net servers with available routes to user and compares their respective off-net ISP charges to evaluate the cheapest off-net server (SRV$_{off}$) to be used in Algorithm 1 for bandwidth cost-efficient user assignment. The greedy heuristic needs to consider the possible routes the traffic can take from the various servers' ISP location to user ISP location. We introduce the notion of Node Relation (NR) map, maintained by CDN, which captures the available routes from servers' ISP to users' ISP and their associated ISP charges. Using the NR map, the greedy heuristic compare the respective off-net ISP charges associated with the available routes from the available off-net servers. Thereafter the choice of cheapest off-net server is driven by the ordering of ISP charging functions. Since customer off-net charging functions are the cheapest, most off-net users are expected to be assigned to off-net servers over customer IP routes.

Algorithm 2 Extended CORE-OMN-Greedy User Assignment

Cheapest-Off-Net-Server
OffNet-SRVs = list of off-net servers
pUser = user to be assigned consuming **b** bandwidth
NR(ISP$_i$, ISP$_j$) = ISP$_i$ off-net traffic charging function
for BGP route from ISP$_i$ to ISP$_j$
for all SRV$_{k,j}$ $\in$ OffNet-SRVs **do**
rel = NR(SRV$_k$·ISP$_j$, pUser·ISP)
Cost$_k$ = C$_j^{rel}$(b)
Find SRV$_{off}$ = SRV with **minimum off-net Cost**
return SRV$_{off}$

C. CORE-OMN User Assignment Problem with Delay Constraints

CORE-OMN user assignment problem focuses on minimizing bandwidth cost alone which can negatively affect the network performance of users and therefore we introduce the notion of delay constraint to the CORE-OMN user assignment problem. The bandwidth cost of group and network performance of users are orthogonal metrics of performance which can not be optimized simultaneously. However the user assignment policy impacts both bandwidth cost and user delay. In order to account for user delay while trying to minimize bandwidth cost we supplement the CORE-OMN user assignment problems (both general and extended) with delay constraint as follows: for each user U$_i$ assigned to SRV$_X$, $dist_{i,X} \leq dist_{i,M}$, where SRV$_M$ is the top Mth server offering least delay to the user. We have found the best trade-off between bandwidth cost and user delay is achieved when user is assigned to one of its top 5 servers offering the least delay.

Algorithm 3 CORE-OMN-Delay constraint User Assignment

Delay Compliant Server Selection
SRVList = original list of servers with available bandwidth
pUser = user U$_i$ to be assigned consuming **b** bandwidth
path-dist$_{x,y}$ = distance between application entities x and y
for all SRV$_k$ $\in$ SRVList **do**
dist$_{i,k}$ = path-dist$_{i,k}$ + path-dist$_{k,root}$
Order SRVList by dist$_{i,k}$
Delay-Compliant-SRVList = Top M servers in SRVList
return Delay-Compliant-SRVList

CORE-OMN-Delay Heuristic: CORE-OMN-Delay exploits the fact that network performance of users is composed of overlay tree delay and last-hop delay in evaluating the delay constraint for each user assignment. Commercial CDNs adopt the nearest available policy which focuses on minimizing last-hop delay and thereafter the distribution mechanism attempts to optimize the overlay tree delay. But the user assignment policy impacts both the overlay tree delay and last-hop delay experienced by the user. For instance, the user assignment policy can assign user to a server more closer to root or more closer to leaf which will result in different overall delay for the user. Therefore imposing delay constraint based on the overall delay from root to user presents a more holistic approach at ensuring better network performance for user.

CORE-OMN-Delay presented in Algorithm 3 outlines the selection of servers which are delay compliant for any given user. The heuristic initially evaluates the top M servers which offer the least overall delay for any given user. Thereafter the user is assigned to the cheapest server amongst the top M servers which are delay compliant by using Algorithm 1 for general and Algorithm 2 in conjunction with Algorithm 1 for extended CORE-OMN user assignment problem.

V. CORE-OMN PROTOCOL

Overlay multicast protocols are involved in (1) assigning users to servers (2) organizing participating servers of the same multicast group into a dissemination tree to deliver content from root to end users and (3) maintaining the overlay tree as group membership changes. CORE-OMN can choose from the various user assignment heuristics based on the underlying ISP charging functions and the desired user performance. Overlay tree construction can be performed using any of the following protocols: OMNI, HMTP, NICE and AMCast. We consider the overlay tree maintenance problem from the point of maintaining low bandwidth cost as group membership changes.

A. CORE-OMN User Movement Problem

Multicast group membership changes as users join and leave. CORE-OMN user assignment ensures bandwidth cost-efficiency for any joining user. However when a user leaves the group a deficiency is created underneath its joining servers which can be fulfilled by any other user. CORE-OMN user movement exploits such opportunity by moving users from costly server locations to any cheap server location where user deficiency occurs to reduce the overall bandwidth cost.

The CORE-OMN user movement problem is stated formally as: Given user deficiency at SRV$_{def}$ find user U$_{i,j}$ under SRV$_{curr}$ for movement which maximizes reduction in bandwidth cost $C^{old}_{curr}(b) - C^{new}_{def}(b)$ under the delay constraint for U$_{i,j}$ and where old, $NR(ISP_{curr}, ISP_j)$, and new, $NR(ISP_{def}, ISP_j)$, captures the old and new IP routes' respective ISP charging functions.

CORE-OMN Distributed User Movement: CORE-OMN Distributed user movement considers the type of bandwidth consumed before and after moving the user to maximize the reduction in bandwidth cost through user movement. The scheme exploits the fact that changing bandwidth consumption of a user from off-net to on-net presents maximum reduction in bandwidth cost. And similarly changing bandwidth consumption of a user from provider off-net to peer or customer off-net and from peer off-net to customer off-net presents the next best alternative for maximum reduction in bandwidth cost. However the scheme also needs to consider the delay constraint associated with user re-assignment to avoid moving users to servers offering unacceptable delay performance.

Algorithm 4 CORE-OMN Distributed User Movement

```
Candidate-User-Move(from ISP_loc, to ISP_def)
  for all U_{i,loc} in ISP_loc and assigned to SRV_{k,curr} do
    old = NR(SRV_k·ISP_curr, ISP_loc)
    new = NR(ISP_def, ISP_loc)
    opportunity-cost = C^old_curr(b) - C^new_def(b)
return U_{i,loc} with maximum opportunity-cost
Search for User Movement
SRV_def within ISP_def with user deficiency
user = Candidate-User-Move(ISP_def, ISP_def)
if user == null do
  for all ISP_loc 1-hop from ISP_def do
    user = Candidate-User-Move(ISP_loc, ISP_def)
  Choose user with maximum opportunity-cost
```

CORE-OMN Distributed user movement presented in Algorithm 4 searches a user for movement which maximizes reduction in bandwidth cost without violating the delay constraint for the moved user. In order to comply with delay constraint the search for user movement is made in increasing scope of distance from ISP$_{def}$ where server SRV$_{def}$ faces the user deficiency. Initially users within ISP$_{def}$ being served underneath any off-net servers are candidates for movement since the change in bandwidth consumption is from off-net to on-net. However if no such user is found then off-net users of nearby ISPs are considered for movement in an attempt to change bandwidth consumption for a user from provider to peer or peer to customer off-net and in that order. The scope of the search is limited to single-hop neighbors of ISP$_{def}$ to satisfy the delay constraint criteria for any user movement and limit the run-time complexity.

VI. Evaluation

In this section we compare user assignment policies on various performance metrics: bandwidth cost, network performance for users and inter-ISP traffic generated by CDN servers and evaluate effectiveness of CORE-OMN user movement policy in reducing bandwidth cost as users join and leave group.

In CDN, servers and users are present within ASes, where server deployment is known and users in any AS can request content. To simulate CDN, AS-level topology [3] providing inter-AS link connectivity is used as the underlying network where servers and users are attached to various AS locations. We use publicly available BGP routing tables from Route-Views [5] and infer the type of BGP policy compliant path taken between any pair of ISPs using Gao's algorithm [14] to construct NR Map. Users are distributed amongst ASes following group membership studies [22], [16] that report the existence of spatial properties such as clustering and diversity in user population of multicast groups. Clustering points out the skew in user population while diversity points out the large number of distinct locations where popular groups are accessed. To simulate both spatial properties we distribute users amongst the physical AS locations following Zipf distribution.

Previous works [15], [27], [13] have reported different types of ISP charging models composed of a ISP charging function, c(x) where x is the charging volume, and ways to determine the charging volume x. In [27], [13] ISPs are reported to charge their customers over the total volume of traffic generated, using the concave charging function $c(r) = (\alpha - \beta \cdot \ln r) \cdot r$, where c is the monthly fee, r is the charging volume in Mbps, and α and β are two parameters. In [15] ISPs are reported to charge their customers over the 95th-percentile of the traffic volume using complex step-wise increasing charging functions. We generate the ISP charging functions for on-net and off-net traffic and for both charging models while maintaining the general trend of charging functions following economy of scale.

A. Analyze Extended CORE-OMN User Assignment Problem

The performance metrics are affected by these variables: ISP locality of servers and users, latency matrix, server traffic load and user traffic demand; information that is proprietary to commercial CDNs. Therefore CORE-OMN protocol with extended CORE-OMN-Greedy and Delay is compared against OMNI protocol with nearest-available, for groups of increasing

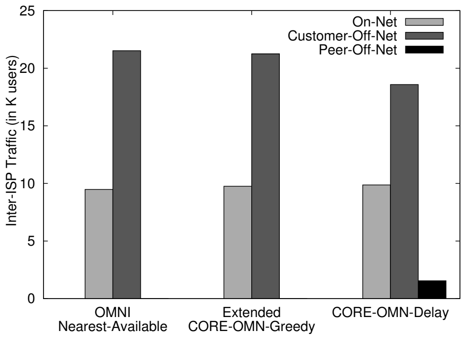

Fig. 2. Each policy generates nearly same on-net and cust. off-net traffic with servers in Top ASes

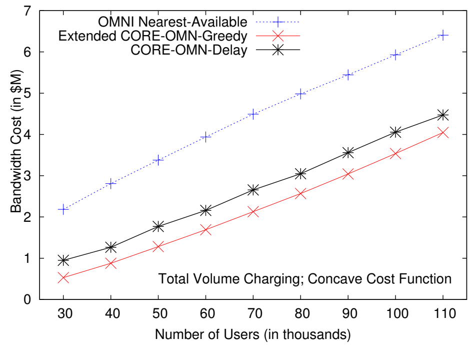

Fig. 3. CORE-OMN-Greedy present least bandwidth cost by assigning users to cheap servers

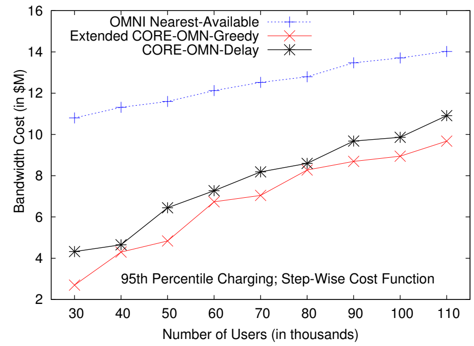

Fig. 4. CORE-OMN-Delay sacrifices on bandwidth cost to improve network performance of users

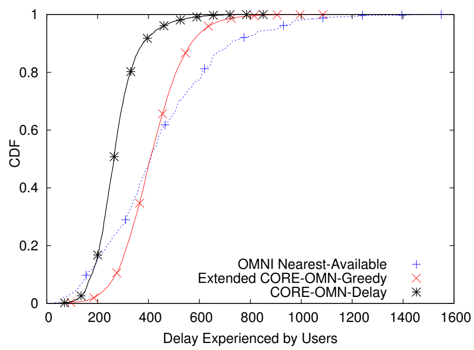

Fig. 5. End-to-End delay constraint improves network performance

sizes where servers are assigned fixed bandwidth and users consume fixed bandwidth, in specific scenarios that realistically capture aforementioned variables.

1) Servers deployed in Top AS locations: With servers deployed in top AS locations of mostly Tier-1 and Tier-2 networks with maximum degree, the user assignment policies mostly generate on-net and customer off-net traffic as shown in Figure 2. On-net server is preferred by CORE-OMN user assignment policies since on-net bandwidth is cheapest due to on-net charging function being cheapest and preferred by nearest-available policy since on-net server is nearest. It explains the near same level of on-net traffic generated by the various user assignment policies. But due to Zipf user distribution, on-net servers in certain locations get saturated forcing other users in those locations to be assigned to off-net servers. As servers are deployed primarily in provider AS locations almost all such off-net users are served over customer paths explaining the elevated levels of customer off-net traffic.

CORE-OMN user assignment policies significantly reduces bandwidth cost as compared to nearest-available policy for both ISP charging models as shown in Figure 3 and Figure 4. The trend for bandwidth cost is same for both ISP charging models and therefore only representative result is shown in other scenarios. Extended CORE-OMN-Greedy presents least bandwidth cost by assigning off-net users to globally cheapest customer off-net servers and saturating them to exploit the concave nature of ISP charging functions. CORE-OMN-Delay assigns off-net users to cheap delay compliant off-net servers, which are not always globally cheapest off-net servers but still offer significant bandwidth cost saving in comparison to nearest-available policy. The nearest-available policy assigns users to nearest available customer off-net server based on user locality that may not necessarily be the cheapest, thereby causing significant increase in bandwidth cost.

CORE-OMN user assignment policies present improved network performance for most users when compared against the nearest-available policy in Figure 5. CORE-OMN-Delay assigns users to off-net servers that are already part of the dissemination tree and therefore offering much lower tree delay, which improves the overall network performance for the users. For certain users such delay compliant and cheap off-net servers are available over peer paths which explains the peer off-net traffic generated by CORE-OMN-Delay in Figure 2. The nearest available server offers minimum last-hop delay which improves network performance for certain percentage of users but the tree delay remains sub-optimal causing an overall increase in delay for every user. CORE-OMN-Greedy by optimizing bandwidth cost assigns users to off-net servers that are cheap but offer greater tree delay thereby sacrificing network performance for certain users.

2) Servers deployed in Tier-1 and Tier-2 Networks: Servers are deployed in every Tier-1 and Tier-2 network to move content closer to users and user distribution is focused on ISPs offering diversity in inter-ISP link connectivity to provide choice in terms of type of inter-ISP traffic generated at CDN servers by user assignment policies as shown in Figure 6. As before, each user assignment policy assigns nearly 30% of user population to on-net servers implying that a certain percentage of user population is bound to be served by nearest available which is also cheapest available on-net server. But there is variety in off-net traffic since same user can be served over different IP routes.

CORE-OMN user assignment policies by regulating off-net traffic in accordance with underlying ISPs' routing preference

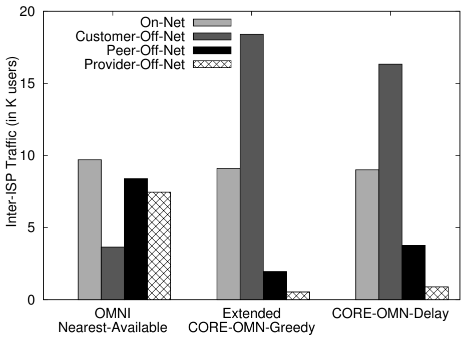

Fig. 6. CORE-OMN shifts provider and peer to cust. off-net traffic with servers in Tier-1,2 ASes

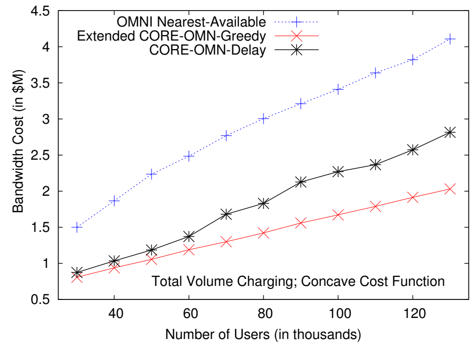

Fig. 7. CORE-OMN reduce bandwidth cost by assigning users to servers over cheap IP routes

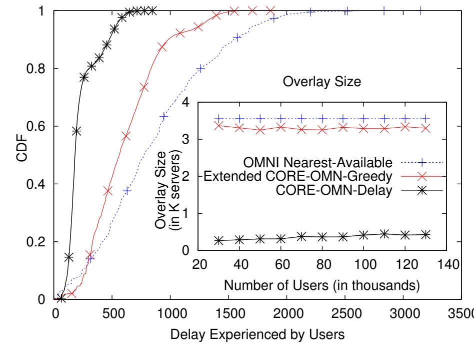

Fig. 8. CORE-OMN-Delay reduces delay by assigning majority of users to few off-net servers

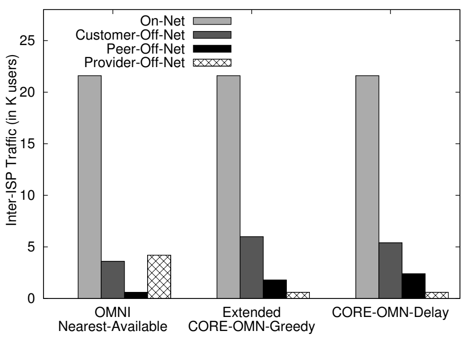

Fig. 9. Servers and users within same ASes

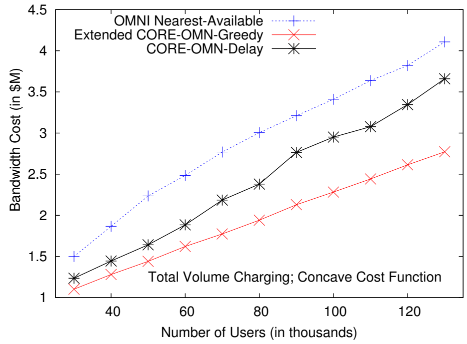

Fig. 10. Reduced ISP cost saving by On-Net SRV

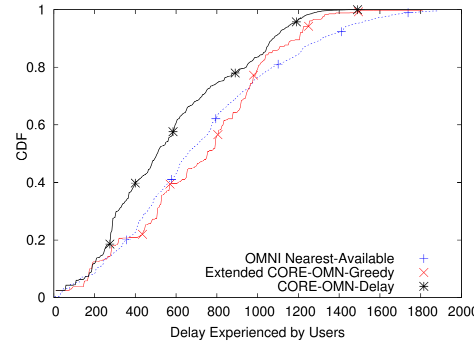

Fig. 11. CORE-OMN-Delay improves user delay

is able to significantly reduce bandwidth cost for CDN when compared to nearest-available policy as shown in Figure 7. CORE-OMN user assignment policies assign off-net users predominantly over customer paths offering cheapest bandwidth costs and only assign users over provider and peer paths when on-net servers in Tier-1 networks are saturated which forces other users in those networks to be assigned to off-net servers in their customer and peer networks. The nearest-available policy assigns users based on their locality which results in random inter-ISP traffic generated over provider, peer and customer paths as seen in Figure 6. CORE-OMN-Greedy minimizes bandwidth cost by assigning most off-net users to few globally cheapest customer off-net servers. The nearest-available policy generates more provider and peer than customer off-net traffic at server sites which adversely affects bandwidth cost. CORE-OMN-Delay by assigning users to cheap delay compliant off-net servers still offers significant bandwidth cost savings.

CORE-OMN user assignment policies significantly improves network performance of users when compared against nearest-available policy as shown in Figure 8, by offering lower tree delay which dominates user delay performance. CORE-OMN-Greedy assigns users to few off-net servers which forms a core dissemination tree which is small in size and therefore lowers tree delay for most users. The nearest available server minimizes the last-hop and thereafter OMNI attempts to minimize the average delay for the users. Since nearest available server joins dissemination tree usually farther away from root it offers significantly higher tree delay. The nearest-available policy presents sub-optimal tree delay to users which increases the overall delay experienced by users. CORE-OMN-Delay by assigning users predominantly to off-net servers which are already part of the dissemination tree is able to reduce the size of overlay tree as seen in Figure 8 which improves the overall network performance for users.

3) Servers deployed to increase On-Net traffic: Servers and users are mostly deployed in same AS location, which increase on-net traffic as shown in Figure 9. Each user assignment policy generates nearly 70% on-net traffic since on-net servers are preferred by each policy for different operational reasons stated earlier. As majority of traffic generated is on-net for which ISP cost is same for each policy, the difference in ISP cost is reduced as shown in Figure 10. For remaining users, CORE-OMN-Greedy reduces overall ISP cost by selecting cheaper IP routes for user assignment and CORE-OMN-Delay by selecting delay compliant off-net servers provides the best network performance.

B. Analyze General CORE-OMN User Assignment Problem

Even large-scale CDNs can deploy servers only in a limited number of AS locations and therefore servers are preferably

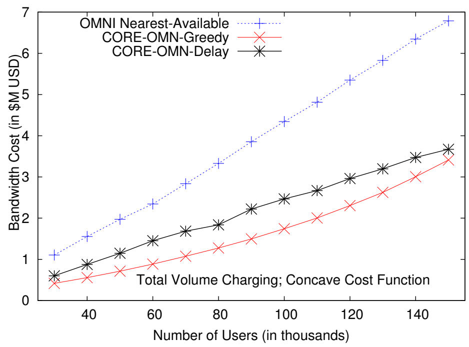

Fig. 12. CORE-OMN-Greedy presents least bandwidth cost with servers in every AS

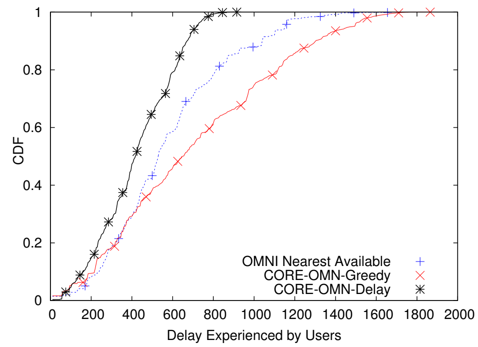

Fig. 13. CORE-OMN-Delay improves user delay by assigning it to delay compliant servers

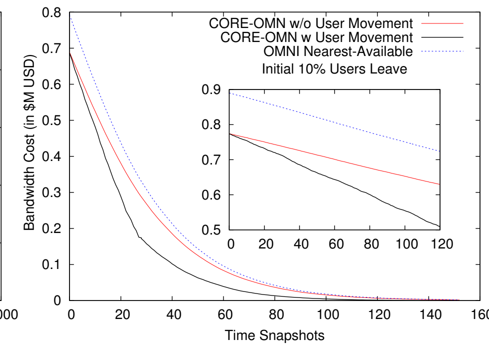

Fig. 14. CORE-OMN with User Movement shows steady drop in ISP bandwidth costs

deployed in ASes with majority of user requests i.e. hot-spot locations. But many such AS locations are customer networks or small regional providers and therefore in this scenario servers are deployed in AS locations where off-net traffic can only default through provider links and users are also distributed amongst these AS locations.

CORE-OMN-Greedy presents lowest bandwidth cost as shown in Figure 12 while CORE-OMN-Delay presents best network performance as shown in Figure 13 when compared to the nearest-available policy. As before, Zipf user distribution causes saturation of on-net servers where other users are then allocated to off-net servers. CORE-OMN-Greedy assigns users to globally cheapest off-net servers to minimize bandwidth cost but sacrifices network performance of users. CORE-OMN-Delay can either assign user to nearest on-net server offering higher tree delay or off-net server offering lower tree delay. CORE-OMN-Delay assigns user to cheapest delay compliant off-net server which presents enough bandwidth savings as well as significant improvement in network performance of user. The nearest available server can either be on-net or off-net which may or may not offer cheap bandwidth cost or better overall delay. And therefore nearest-available policy is unable to control bandwidth cost or provide better network performance.

C. Summary

CORE-OMN policies minimize bandwidth cost for CDN by redistributing traffic over cheaper IP routes. The bandwidth cost depends upon ISP charging function and therefore associating ISP charging functions with different types of traffic allows bandwidth cost to depend upon the type of traffic generated at server sites. On-net traffic is the cheapest due to on-net ISP charging function being the cheapest and therefore always preferred by CORE-OMN policies. But since servers are not deployed in every AS location, off-net traffic at server sites is unavoidable. CORE-OMN policies are able to redistribute off-net traffic from costly provider and peer IP routes to cheaper customer IP routes due to the ordering of the associated provider, peer and customer off-net ISP charging functions. CORE-OMN policies reduce bandwidth cost by redistributing off-net traffic irrespective of the type of ISP charging model adopted as long as the underlying routing preferences of ISPs are communicated through the ISP charging functions.

Incentives are available for both CDN and ISPs to participate in such a price sharing mechanism. CDNs can reduce their operational bandwidth cost which can be significant in case CDN traffic keeps getting assigned to costly IP routes. ISPs by providing the necessary economic incentive can motivate the redistribution of CDN traffic over cheaper IP routes which will reduces ISPs own cost of carrying the CDN traffic.

D. Analyze CORE-OMN Distributed User Movement

CORE-OMN protocol with user movement is compared against CORE-OMN protocol without user movement and OMNI protocol, on bandwidth cost as group membership changes. To simulate changes in group membership initially 30K users are allowed to join servers in top Tier-1 and Tier-2 networks using the various user assignment policies and thereafter users are randomly chosen to leave the group. Snapshots of bandwidth cost are taken at regular intervals for comparison.

CORE-OMN distributed user movement presents most reduction in bandwidth cost as initial 10% users leave, as shown in Figure 14, when snapshots are taken after every 50 users leave. Initially as users leave from saturated on-net server locations other users in those locations which had been assigned to off-net servers get chance for movement underneath their on-net servers. Moving users from off-net server locations to on-net server locations produces maximum reduction in bandwidth cost. Therefore as more users leave from saturated locations, costly off-net bandwidth consumption is replaced by cheaper on-net bandwidth consumption which produces the sharp drop in bandwidth cost. CORE-OMN without user movement and OMNI do not exploit such opportunities and therefore the drop in bandwidth cost is only due to user leaving.

CORE-OMN distributed user movement reassigns users to cheaper server locations with user deficiency as long as it reduces bandwidth cost and the delay constraint is met for moved user. Figure 14 presents the complete trend of drop in bandwidth cost as all users leave the multicast group and when snapshots are taken after every 5% change in group membership. Initially as users leave saturated server locations, more opportunities exist for moving users to servers which reduce bandwidth cost while meeting the delay constraint. However at certain stage most users have already been moved underneath their respective on-net servers and thereafter less viable opportunities exist for user movement. So thereafter drop in bandwidth cost is only due to user leave which is marginal since these users have already been moved to their respective on-net server locations. CORE-OMN without user movement and OMNI again drop bandwidth cost only due to user leave.

VII. Related Work

In recent times Internet-scale dissemination of video content is being achieved through CDNs where dedicated servers act as proxies facilitating the multicast groups. The advantage of CDN is end users send or receive only one copy of data packets during session, and the work of duplicating packets is shifted from data sources to servers.

CDNs select the best route based on global information about link delays which may violate business agreements about traffic routing between ISPs [10]. CDNs enable service and content distribution costs to be shared amongst multiple providers but since the traffic patterns also determine money flows between providers, CDNs may also influence commercial relationships on the Internet. Network operators are seeking ways to mitigate the side-effect of CDN routing on the ISPs. Jiang et al. [17] study a joint design system where ISPs and CDN cooperate to achieve both ISP's traffic routing and CDN's user performance goals. But the optimal solution is achieved only when CDN gains complete control of routing for its content traffic with complete visibility into ISP's network i.e. routing decisions on OSPF weights, real-time link latency, traffic matrix etc. while the ISP solves routing problem only for background traffic. In contrast, CORE-OMN aligns CDN and ISPs routing preferences by sharing only pricing information of routes.

Overlay multicast protocols have traditionally focused upon application level performance objectives when deploying multicast groups. *AMCast* [21], *OMNI* [8] and ROMaN [18] are state-of-art OMN protocols used to organize servers into overlay tree facilitating the multicast service for end-users. *AMCast* protocol attempts to optimize bandwidth usage at the servers to maximize the number of groups. *OMNI* protocol attempts to minimize the average delay experienced by users in a multicast tree through local transformations whenever any change occurs in network conditions and user memberships. ROMAN protocol attempts to minimize ISP charge of multicast group by redirecting user request to the cheapest available server without considering underlying ISP transmission costs. Only CORE-OMN protocol attempts to align CDN and ISP routing preferences for routing content traffic with minimal information sharing.

VIII. Conclusions

ISPs, by charging separately for on-net and different types of off-net traffic, are able to recover the actual cost of delivering CDN traffic as well as provide necessary economic incentives for CDNs to consider ISPs' routing preferences. In response, CORE-OMN protocol minimizes CDN payment to ISPs by choosing cheaper IP routes for user assignment. CORE-OMN-Delay by assigning users to delay compliant servers on cheap IP routes provides best trade-off *i.e.* good network performance with significant savings in bandwidth cost. CORE-OMN-Greedy offers least bandwidth cost but at the expense of network performance of users. The nearest-available server minimizes last-hop delay but higher tree-delay increases the overall user delay. As for off-net users, when nearest-available server is on costly IP route it will increase operational cost for ISPs and CDN. For future work we plan to investigate a distributed algorithm for CORE-OMN user assignment problem. We also intend to apply this economically incentivized cooperation scheme to non-multicast services offered by CDNs.

References

[1] 2010 NCAA March Madness Live. http://mmod.ncaa.com/.

[2] Akamai - Powering a better Internet. http://www.akamai.com/.

[3] Internet Topology Collection. http://irl.cs.ucla.edu/topology.

[4] Limelight Networks. http://www.limelightnetworks.com/.

[5] RouteViews Project. http://www.routeview.org.

[6] The Great Obama Traffic Flood. http://asert.arbornetworks.com/2009/01/the-great-obama-traffic-flood/.

[7] Cisco Visual Networking Index: Forecast and Methodology, June 2009. http://www.cisco.com/US/white_paper_c11-481360.pdf.

[8] S. Banerjee, C. Kommareddy, K. Kar, B. Bhattacharjee, and S. Khuller. Construction of an efficient overlay multicast infrastructure for real-time applications. In *Proc. of IEEE INFOCOM*, 2003.

[9] Y.-H. Chu, S. G. Rao, and H. Zhang. A case for end system multicast. In *Proc. of ACM SIGMETRICS*, 2000.

[10] D. Clark, B. Lehr, S. Bauer, P. Faratin, R. Sami, and J. Wroclawski. The Growth of Internet Overlay Networks: Implications for Architecture, Industry Structure and Policy. In *TPRC*, 2005.

[11] R. D. Day. Meta Content Delivery Network. *US Patent 7185052*, 2007.

[12] J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, and B. Weihl. Globally Distributed Content Delivery. *IEEE Internet Computing*, 2002.

[13] P. M. Ferreira. Implications of decreasing bandwidth price on allocating traffic between transit and peering aggreements. *MS Thesis, MIT*, 2002.

[14] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9:733–745, 2000.

[15] D. K. Goldenberg, L. Qiu, H. Xie, Y. R. Yang, and Y. Zhang. Optimizing cost and performance for multihoming. In *Proc. of SIGCOMM*, 2004.

[16] J. hong Cui, M. Faloutsos, D. Maggiorini, M. Gerla, and K. Boussetta. Measuring and modelling the group membership in the Internet. In *Proc. of IMC*, 2003.

[17] W. Jiang, R. Zhang-Shen, J. Rexford, and M. Chiang. Cooperative content distribution and traffic engineering in an ISP network. In *Proc. of ACM SIGMETRICS*, 2009.

[18] V. Khare and B. Zhang. Towards Economically Viable Infrastructure-based Overlay Multicast Network. In *Proc. of IEEE INFOCOM*, 2009.

[19] G. Pallis and A. Vakali. Insight and Perspective for Content Delivery Networks. In *Communications of the ACM*, 2006.

[20] D. M. Shaw. Streaming media subscription mechanism for a content delivery network. *US Patent 6751673*, 2004.

[21] S. Shi and J. Turner. Routing in Overlay Multicast Networks. In *Proc. of IEEE INFOCOM*, 2002.

[22] K. Sripanidkulchai, B. Maggs, , and H. Zhang. An Analysis of Live Streaming Workloads on the Internet. In *Proc. of IMC*, 2004.

[23] L. W. Vivek, L. Wang, V. Pai, and L. Peterson. The Effectiveness of Request Redirection on CDN Robustness. In *Proc. of 5th OSDI*, 2002.

[24] F. Wang and L. Gao. On Inferring and Characterizing Internet Routing Policies. In *Proc. of ACM IMC*, 2003.

[25] B. Zhang, S. Jamin, and L. Zhang. Host Multicast: A Framework for Delivering Multicast To End Users. In *Proc. of IEEE INFOCOM*, 2002.

[26] L. Zhang, S. Floyd, and V. Jacobsen. Adaptive Web Caching. In *Proc. of the NLANR Web Cache Workshop*, 1997.

[27] Y. Zhu, C. Dovrolis, and M. Ammar. Combining multihoming with overlay routing. In *Proc. of IEEE INFOCOM*, 2007.